Supplementary Materials for

The Developmental Morphospace of Wildfire

Ty Tuff, Ph.D.

Environmental Data Science Innovation & Impact Lab (ESIIL), Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado Boulder, Boulder, Colorado, USA.
ORCID: 0000-0001-5249-5197
Correspondence: ty.tuff@colorado.edu

Contents

Supplementary Methods

S1. Source observations, date audit, and analysis populations

Fire VASE used FIRED event and daily observation products derived from MODIS burned area [1, 2, 3]. Event records supplied identifiers, catalog area, duration, dates, and final geometry. Daily records supplied event identifier, observation date, newly mapped area, cumulative mapped area, and daily geometry. Every retained increment was joined to the source daily GeoPackage by event identifier and date and required to agree within 10^{-9} km^2. Catalog areas were checked against the event GeoPackage. The raw daily source contains 627,033 rows, whereas the analysis-ready data lake contains 626,102. The repository exclusion table identifies the 931 unmatched source rows and the events to which they belong.

The archive contains 278,569 events dated from 2 November 2000 through 1 May 2021. No retained event has a missing or duplicated observation date, a nonfinite or negative increment, an observed zero increment, or a zero reconstructed total. Such records were excluded rather than converted to zero growth. Reconstructed totals and catalog totals agree to within 6.5×10^{-16} in relative magnitude in the cached events.

Observation adjacency is not equivalent to calendar-day adjacency. Of 347,533 adjacent observed pairs, 196,611 are exactly one day apart, 81,940 are two days apart, and 68,982 span more than two days. An unobserved date was never filled with a zero-growth row. The full archive includes 161,073 one-observation events and 47,950 two-observation events. Among the 69,546 events with at least three observations, 59,300 contain at least one gap. The remaining 10,246 have no missing calendar dates between their first and last observation and form the primary developmental population (Table S1). Thus "primary" does not mean that an event merely contains one three-day run; its complete retained history is gap free.

Three additional populations serve different analytical purposes. The primary event-weather cohort contains 9,212 primary fires with complete values for every requested weather variable. The archive-wide weather-complete set contains 237,235 events, but it is used only to describe weather-data availability and is not the primary morphology-model cohort. The state-model cohort contains 87,944 exact subsequent-calendar-day transitions from 31,700 fires. The matching analysis uses the same 9,212 fires as the primary event-weather analysis.

S2. Normalized VASE construction and developmental summaries

For event i with dated nonnegative increments g_{ij}, reconstructed burned area was

$$S_i = \sum_{j=1}^{n_i} g_{ij}, \tag{S1}$$

and normalized growth allocation was $p_{ij} = g_{ij}/S_i$. Cumulative reconstructed area through time t was $A_i(t)$. VASE width was

$$w_i(t) = \sqrt{\frac{A_i(t)}{A_i(T_i)}}. \tag{S2}$$

Positive-area histories therefore terminate at width one. The square root maps relative area to a radius-like dimension; width does not describe the geographic perimeter, displacement, or spread direction.

Within the primary consecutive cohort, every dated increment occupied one equal-width interval in relative developmental time. Cumulative mass was linearly interpolated at 21 fixed edges and differenced to produce 20 nonnegative allocation bins. The bins sum to one and conserve reconstructed area allocation. Interpolation does not create additional daily observations. Gappy histories enter only designated observation-time sensitivities; missing dates are not reconstructed.

The visual and quantitative representations use the same history in different forms. A VASE displays cumulative allocation through its radius, whereas the primary PCA uses the 20 incremental allocation bins. Consequently, similarity between rendered silhouettes is not itself the distance metric used in the morphospace.

Scalar summaries were calculated after the primary coordinate system was fitted. Front-loaded fraction is allocation in the first half of relative development; late-growth fraction is allocation in the final quarter. Peak timing is the relative position of the largest observed daily increment. Terminal taper is final observed growth divided by the observed daily peak. Shannon entropy is

$$H_i = -\sum_j p_{ij} \log p_{ij}, \tag{S3}$$

with zero-probability terms omitted; normalized entropy is $H_i/\log n_i$ for $n_i > 1$ [4]. A positive one-observation history has entropy zero but no within-event developmental trajectory. Normalized first and second differences summarize changes between adjacent standardized allocation bins; they are not interpreted as physical velocity or acceleration.

The corrected observed daily peak is $\max_j(g_{ij})$. It replaces the pilot field that divided catalog area by catalog duration; 108,214 event values differ between those definitions. The corrected quantity is a dated daily maximum, not an instantaneous or hourly spread rate.

S3. Shape-only PCA and interpretive landmarks

The primary PCA uses only allocation bins 00 through 19. Final area, catalog duration, observation count, observed daily peak, mean growth, slenderness, scalar developmental traits, and all weather variables are excluded. Each training column was centered on its mean and divided by its training standard deviation; near-constant columns received a scale factor of one. PCA was fitted by eigendecomposition of the standardized covariance matrix [5]. Because the nonnegative bins sum to one, they form a composition and contain a linear constraint. The primary analysis treats them as a standardized Euclidean coordinate representation. A deterministic sign convention makes the stored loadings reproducible. Component scores retain their original variance and are not whitened.

The compositional sensitivity instead used the Hellinger transformation $h_{ij} = \sqrt{p_{ij}}$. Allocation rows were required to be finite, nonnegative, and sum to one within 10^{-8}; the transformed rows were required to have squared Euclidean norm one within the same tolerance. Hellinger columns were mean centered without variance scaling, so their Euclidean distances retained Hellinger geometry up to a constant. PCA was fitted to the same 10,246 events by deterministic singular value decomposition. Components were assigned with the Hungarian algorithm using maximum absolute score correlation and were sign oriented to positive aligned correlation. Raw allocation contrasts at score tails aided temporal interpretation; loading magnitudes were not compared directly across transformations because their units differ.

Interpretive landmarks were assigned after coordinate construction. Rules first identify invalid, one-observation, two-observation, and gappy histories. Among eligible consecutive histories, multiple detected pulses require at least two local maxima with prominence at least 20% of the observed daily peak. Series endpoints are eligible after zero padding, and a flat plateau counts once. Reactivation requires growth to return above 25% of the peak after at least two consecutive observed subthreshold days. A late peak has normalized peak-bin midpoint at least 2/3; a front-loaded taper has at least 75% of area allocated by relative time 0.5 and final-to-peak growth at most 0.35. Remaining eligible histories are described as distributed growth. These declared rules provide landmarks; they are not fitted clusters or evidence for or against natural fire classes.

S4. Stability and temporal-order reference analyses

The unchanged 20-bin PCA was refitted at minimum consecutive-observation thresholds of 2, 3, 5, and 7. It was also refitted across bootstrap samples, individual years, and broad regions. We then projected the same seeded set of 1,000 primary fires into every fitted space, providing common reference points for comparison. Axis matching accounted for arbitrary component order and sign. Global stability was measured from correlations between axis scores, similarity of loading vectors, overlap between the five-dimensional subspaces, Procrustes similarity, and rank correlation of pairwise distances. Local stability was measured from overlap among each fire's 15 nearest neighbors and among the most extreme 2% of axis exemplars. The same fixed 1,000-event anchor sample and metrics were used to compare primary and Hellinger geometry. These tests compare alternative fits of the same dataset; they are not independent external validation.

The temporal-order comparison used a recorded sample of 4,000 primary fires. Shuffling within each fire preserved its observation count, reconstructed total, relative-time support, and complete set of increments; only their order changed. Symmetric Dirichlet allocations with concentration parameters of 1 and 10 preserved observation count and total area but supplied alternative positive growth allocations. We compared 100 replicates with the fixed observed sample. Empirical tail probabilities were calculated as $(1 + \text{exceedances})/(1 + \text{replicates})$ and were descriptive rather than adjusted for multiple comparisons. These simulated references test temporal asymmetry and the behavior of the representation under specified nulls; they are not additional observational evidence or tests of a universal developmental program.

S5. Endpoint projections

Endpoint attributes excluded from PCA were projected onto the fitted scores using Spearman correlation. Attributes included final area, duration, observation count, true observed daily peak, and mean growth. Projection asks whether developmental form covaries with event scale or observation length after those variables have been prevented from defining the axes. It does not test statistical or causal independence.

S6. Weather completeness, association, and event-level prediction

Daily gridMET supplied maximum and minimum temperature, VPD, wind speed, precipitation, relative and specific humidity, 100-hour and 1000-hour fuel-moisture indices, energy release component, burning index, reference and potential evapotranspiration, and solar radiation [6]. The fire-danger indices are modeled products, not direct measures of local fuel load. Event histories required complete values for every requested variable on every observed date.

All 237,235 archive-wide complete events are in CONUS; 41,279 Alaska and 55 Hawaii events are incomplete. Complete events have mean duration 3.34 days and mean observation count 2.32, compared with 2.35 days and 1.82 observations among incomplete events. The primary event-weather cohort contains 9,212 fires.

Event-level weather representations include means, quantiles, and fractions exceeding fixed thresholds. The VPD exceedance threshold is 2 kPa, and a wet day has precipitation above 0.1 mm day^{-1}. Because longer records provide more opportunities to observe extremes, these summaries remain sensitive to record length. For adjusted rank associations, ranked weather and ranked developmental responses were each regressed on duration, observation count, area, region, calendar month, and year; the correlation was then calculated between the residuals. Duration and observation count are strongly related and serve only as adjustment variables, not as separately estimated effects. Resampling intervals grouped observations by fire, year, or four broad regions and treated the fitted residuals as fixed.

Ridge models used standardized predictors and penalties 0.01, 1, and 100; penalty 1 was prespecified. Nested sets separately evaluated core means, maximum VPD, comprehensive means, and explicit additions of maximum VPD. All directly compared models share the same complete cohort. The main text displays core means plus maximum VPD because this compact common specification is available for every response on the identical cohort; it is not asserted to be universally optimal. For shape-PC prediction, PCA was refitted within each training fold and held-out scores were centered consistently with the fold-specific target.

Validation held out each broad region, contiguous five-year blocks beginning in 2000, or strict combinations of region and year block. In the strict space–time test, training excluded both the test region and the test year block. Random assignment of fires to folds was retained only as an interpolation diagnostic. Held-out R^2 used a reference mean calculated from the training data, and negative values were retained. Confidence intervals used 200 paired resamples of the existing out-of-fold errors; they do not include uncertainty from refitting the models, remote sensing, or weather exposure.

S7. Exact subsequent-calendar-day state models

State models use only adjacent observations separated by exactly one calendar day and require an observation on the previous calendar day for every comparator. The common cohort contains 87,944 transitions from 31,700 fires. The outcome is $\log[1+g_{t+1}(\mathrm{km}^2)]$. The autoregressive state baseline contains current growth, previous-day growth, cumulative reconstructed area, and elapsed time. Current-growth persistence and day-t weather-only models provide additional reference points.

Day-t newly burned polygons were projected before their centroids were calculated; centroids outside the weather grid were excluded. These fire and weather quantities are retrospectively reconstructed end-of-day covariates. Final-event-centroid exposure was retained only as a sensitivity comparison.

Additive models append day-t weather to the state baseline. Interaction models also include products between weather and state variables. VPD-specific tests add or remove the interactions of VPD with current and previous growth while keeping the observations, validation folds, other weather variables, and penalty unchanged. Paired bootstrap intervals quantify changes in held-out performance using the fixed predictions. Descriptive coefficients on the original measurement scale use standard errors clustered by fire and remain conditional on the fitted design and ridge penalty. Refits by region, year, season, fire area, and training fold assess heterogeneity. Tables of VPD by fire state report both transitions and distinct fires in each quantile cell; tied values can merge cell boundaries. We do not extrapolate a response surface into unsupported portions of the data.

S8. Caliper-constrained matching and conditional references

Matching uses either core event-mean weather variables or the 20 morphology bins, standardized within the 9,212-fire cohort. For each fire, potential partners are its 10 nearest neighbors within the same region, season, duration, and observation-count group. A potential pair is rejected if the larger fire is more than twice the area of the smaller fire or if its root-mean-square standardized distance exceeds 0.5. The remaining pairs are sorted by distance and event identifiers and selected in order without reusing a fire. Differences in the other representation do not influence which pairs are selected.

Candidate coverage is the fraction of fires with at least one acceptable partner among the specified k nearest neighbors; failure to find a match does not prove that no suitable partner exists elsewhere. Sensitivity analyses use $k = 1, 5, 10$, and 20, compare Euclidean/root-mean-square with cityblock/mean-absolute distance, and vary the root-mean-square distance limit from 0.25 to 1. Conditional reference distributions shuffle the other representation within the same adjustment groups, with an additional area class defined by $\lfloor \log_2(\text{area}) \rfloor$. Groups containing only one fire cannot be shuffled.

S9. Computational provenance and verification scope

All results derive from the analysis-ready v0.1 FIRED/gridMET data lake. Missing required inputs cause the workflow to stop, and no synthetic fallback is permitted. The validated repository baseline SHA is `18c923cb0c82bf9f66567b62a3491ac30a28c369`. The analysis configuration is `config/analysis_v2.json` with seed 20260828. Run, publication, validation, and freeze manifests record hashes for the source data, weather cache, code, configuration, and publication outputs.

Automated checks confirmed the scientific invariants, claim matrix, null conservation, manuscript values, publication hashes, figure copies, and deterministic regeneration. All 66 previously frozen outputs were unchanged. The compositional package is additive and records its own input, configuration, output, and test hashes. Its 35 dedicated tests pass. The broader available suite passes 137 tests with two skipped; two legacy modules remain excluded because they import a missing helper unrelated to the compositional analysis.

Supplementary Results

S10. Population support and primary shape geometry

The primary consecutive-history criterion retained 10,246 of 278,569 archive events (3.68%). These are gap-free histories with at least three observations, not a weighted sample of the full archive, and no extrapolation to excluded events was attempted. Retention varied from 2.49% in Alaska and 2.61%

in the central region to 4.90% in the West and 5.53% in the intermountain region (Table S8). It also increased strongly with area, from 0.87% among events no larger than 1 km^2 to 18.38% at 10–100 km^2 and 30.48% at 100–1000 km^2. By recorded duration, 27.67% of three-day histories were retained, compared with 4.08% of histories lasting more than five and no more than ten days. The primary population therefore emphasizes larger events and short, fully observed sequences rather than representing the archive uniformly.

Minimum consecutive-observation thresholds retained 27,499 events at two observations, 2,887 at five, and 1,171 at seven (Table S1). The primary 20-bin PCA explained 34.15% of standardized variance on PC1, 28.2% on PC2, and 14.5% on PC3; the first five axes explained 89.44%.

Observation-depth refits supported broad, but not complete, stability (Table S3). With a minimum of seven observations, correlations for the first two axes were 0.985 and 0.997 when the same reference fires were projected into both spaces. The rank correlation between pairwise distances was 0.969. Procrustes similarity remained 0.981, and overlap between the five-dimensional subspaces was 0.894. In contrast, nearest-neighbor overlap was 0.718, overlap among extreme exemplars was 0.606, and variance explained by the first five axes was 74.21%. The first two developmental directions therefore remain recognizable, whereas local membership, extreme examples, and higher-dimensional coverage are less stable.

Across bootstrap refits, five-dimensional subspace overlap had a median of 0.999 and a 2.5th–97.5th percentile range of 0.998–1.000. Refits by region and year also generally retained the leading geometry. Resolution altered the concentration of variance: the first five axes explained 94.18% with 10 bins, 89.44% with the primary 20 bins, and 87.50% with 40 bins; corresponding PC1 values were 37.24%, 34.15%, and 33.12%. Figures S1 and S2 show the detailed axis, loading, resolution, and local-overlap diagnostics.

The Hellinger fit retained similar variance concentration and global organization (Fig. S4; Table S9). PC1 explained 33.79% and the first five axes 88.28%, compared with 34.15% and 89.44% in the primary fit. Aligned PC1 and PC2 score correlations were 0.94197 and 0.93833. Five-dimensional Procrustes similarity was 0.98479, score-subspace overlap was 0.96009, and pairwise-distance Spearman correlation was 0.97629. Local agreement was lower: mean 15-neighbor overlap was 0.70140 and exemplar-tail Jaccard overlap was 0.43027. Baseline and Hellinger PC1 remained strongly related to front loading ($\rho = -0.96834$ and -0.94502), preserving the dominant early-versus-late allocation gradient. The secondary concentrated-versus-distributed contrast remained visible but rotated and weakened. These results support robustness of broad geometry while showing that exact local membership and tail exemplars depend more strongly on the metric.

S11. Temporal ordering and endpoint projections

Within-fire shuffling changed traits that depend on order while preserving the complete set of observed increments (Table S4). Mean first-half allocation was 0.54063 in observed histories and 0.49958 after shuffling. Mean pulse count was 1.24900 observed and 1.41280 shuffled; mean reactivation count was 0.03800 and 0.12031. Entropy was unchanged by construction. These comparisons identify systematic temporal asymmetry relative to the specified within-fire permutation; they do not identify its physical cause.

Five-axis compression did not uniquely distinguish observed histories from all positive-allocation references. In the 4,000-fire sample, observed first-five-axis variance coverage was 0.89582, compared with 0.85545 after shuffling, 0.90158 for Dirichlet concentration 1, and 0.90853 for Dirichlet concentration 10. Low dimensionality alone therefore does not establish biological restriction.

Endpoint projection showed moderate associations with shape. PC1 correlated −0.308 with final area, −0.185 with duration and observation count, and −0.332 with observed daily peak. PC2 correlated 0.259 with final area and 0.256 with duration and observation count. Preventing endpoints from defining the coordinates does not make developmental morphology statistically independent of event scale or length.

S12. Weather selection, adjusted associations, and response-specific prediction

The 9,212-fire primary event-weather population is selected geographically and observationally. It excludes all Alaska and Hawaii events in the available weather cache and favors longer, more frequently observed histories. Accordingly, weather results apply to the complete primary CONUS population rather than every FIRED event.

After adjustment for duration, count, area, region, month, and year, highlighted weather associations were small (Table S5). Mean VPD correlated 0.04989 with front-loaded fraction and −0.06936 with late allocation. Mean precipitation correlated 0.04072 with late allocation. Its raw association with pulse count was 0.07004 but its adjusted association was 0.02492; the four-region conditional interval (−0.00502, 0.06360) crossed zero.

Weather prediction varied by response and validation scheme. Among the tested ridge specifications, the core-means-plus-maximum-VPD set displayed in the main text had region-held-out R^2 of 0.0215 for front loading, 0.0383 for late allocation, 0.0005 for peak timing, 0.0191 for pulse count, 0.0220 for reactivation, 0.0744 for normalized entropy, 0.0803 for terminal taper, 0.1755 for normalized first difference, and 0.1606 for normalized second difference. Strict crossed space–time values were similar or slightly lower. Fold-specific shape PCs were weakly recovered. These results characterize the declared summaries and ridge models; they do not rule out information recoverable by other exposure definitions or nonlinear model classes. No cross-response median is interpreted as a general weather-performance result.

S13. State baseline and qualified VPD increment

As expected for a sequential growth response, the autoregressive model provided most exact subsequent-day predictive information (Table S6). Held-out R^2 was 0.45838 under random-fire assignment, 0.45822 under year blocks, 0.44764 under region blocks, and 0.44753 under strict crossed space–time validation. The substantive comparison is the increment above this baseline. Adding day-t weather produced increments of 0.00539, 0.00511, 0.00501, and 0.00494. Adding all interactions between weather and fire state produced increments of 0.01522, 0.01465, 0.01812, and 0.01797.

Across validation schemes, the two VPD interaction terms added 0.006–0.012 R^2 beyond the other weather–state interactions. Under region holdout, the increment was 0.01179, with a fire-bootstrap interval of 0.0097–0.0138. In the pooled ridge model expressed in original units, the coefficients for VPD interacting with current and previous growth were 0.21669 and -0.11798 per kPa times $\log(1 + \mathrm{km}^2)$; their fire-clustered standard errors were 0.0127 and 0.0105. Coefficients varied by region, fire size, and partial edge year. The central 1st–99th percentile VPD range was 0.19–3.48 kPa. These results support a small VPD contribution that depends on current fire state, not a universal response curve.

S14. Matching coverage and null-compatible mismatch

Among 9,212 eligible fires, weather-space matching formed 3,710 unique pairs and morphology-space matching formed 3,145. These correspond to 80.5% and 68.3% of eligible fires being paired. Among the ten nearest eligible neighbors, 91.2% of fires had at least one acceptable weather match and 81.3% had at least one acceptable morphology match. As the root-mean-square distance limit increased from 0.25 to 1, paired coverage increased from 59.8% to 85.7% for weather matching and from 54.0% to 77.9% for morphology matching.

Observed mismatch was compatible with conditional permutation references (Table S7). Among weather matches, the fraction with morphology distance greater than one was 0.49730, compared with null mean 0.50401 and central reference interval 0.49204–0.51834. Among morphology matches, the fraction with weather distance greater than one was 0.39491, compared with null mean 0.40227 and interval 0.39014–0.41750. No excess mismatch was detected under the declared design.

The two main-text examples are independent pairs selected near the median population mismatch. They are controlled candidates for follow-up, not maxima, a reciprocal two-by-two arrangement, or evidence identifying a missing ecological mechanism.

Supplementary Figures

Validated v2 Figure S1

Publication asset `Supplementary_Figure_1.pdf` or `Supplementary_Figure_1.png`.

Open the validated repository figure.

Figure S1: **Complete supplementary validation and sensitivity diagnostics.** (A) Region-held-out R^2 across responses and nested predictor sets. (B) First-five-axis explained variance for alternative feature, observation-support, and resolution choices. (C) Bootstrap-aligned loading stability for PCs 1–5. (D) Weather completeness by broad region. (E) Fraction paired across candidate-neighbor counts and distance metrics. (F) Region-held-out state-model performance across ridge penalties. Complete machine-readable tables accompany every panel.

Validated v2 Figure S2

Publication asset `Supplementary_Figure_2.pdf` or `Supplementary_Figure_2.png`.

Open the validated repository figure.

Figure S2: **Robustness of developmental morphology and limits of external interpretation.** (A) Similarity of pairwise-distance ranks, 15-neighbor sets, and extreme exemplars as the minimum number of consecutive observations increases. (B) Variance explained by the first five axes for observed histories, within-fire shuffles, and two Dirichlet positive-allocation references in the same 4,000-fire sample. (C) Observed and shuffled means for front loading, detected pulses, and reactivation. (D) Spearman correlations of final area, duration, observation count, and observed daily peak with PC1 and PC2. (E) Year-adjusted partial-rank weather associations with intervals obtained by resampling regions. (F) Fraction of fires assigned a unique partner as the matching distance limit changes. Broad developmental gradients are stable, but local structure, compression, environmental associations, and matching coverage require more qualified interpretation.

Validated v2 Figure S3

Publication asset `Supplementary_Figure_3.pdf` or `Supplementary_Figure_3.png`.

Open the validated repository figure.

Figure S3: **Data support and heterogeneity of VPD–state interactions.** (A,B) Numbers of transitions across quantile groups of VPD and current or previous growth; color gives $\log_{10}(\text{transitions}+1)$. (C) Coefficients for interactions between VPD and growth in the pooled sample and in regional refits, expressed in the original units. (D) Increase in held-out R^2 from the VPD interaction terms above the additive state-plus-weather model and above models containing only the other weather–state interactions. Crosses mark cells with fewer than 100 transitions or 30 fires. Tied growth values can merge quantile boundaries. No response surface is extrapolated beyond the observed data.

Validated compositional sensitivity

Add `compositional_sensitivity.pdf` or `compositional_sensitivity.png`.

Open the validated repository figure.

Figure S4: **Compositional-geometry sensitivity of the developmental morphospace.** (A) Primary standardized-Euclidean PCA scores. (B) Hellinger PCA scores after component assignment and sign alignment. (C) Relationships between aligned PC1 and PC2 scores. (D) Differences in raw temporal allocation between upper and lower score tails, used to compare axis interpretation without treating loadings in different transformed units as directly equivalent. (E) Global and local correspondence metrics. (F) Movement of the exemplar fires displayed in the main figures. Broad configuration and the dominant early-versus-late gradient persist under Hellinger geometry, whereas neighborhood membership, extreme exemplars, and the secondary concentrated-versus-distributed contrast are less invariant.

Supplementary Tables

Table S1: **Population definitions and sample sizes.**

Population	Events or rows	Fires	Analytical role
Full FIRED archive	278,569 events; 626,102 observations	278,569	Source population and observation-support audit
One observation	161,073 events	161,073	Observed outcome without internal trajectory
Two observations	47,950 events	47,950	Minimal ordering; 17,253 have two consecutive observations
At least three observations	69,546 events	69,546	Includes 59,300 gappy events
Primary, at least three observations and gap free	10,246 events	10,246	Shape-only PCA and developmental inference
At least five consecutive	2,887 events	2,887	Observation-depth refit
At least seven consecutive	1,171 events	1,171	Strict observation-depth refit
Archive weather complete	237,235 events	237,235	Completeness and selection audit; CONUS only
Primary weather complete	9,212 events	9,212	Event-level weather models and matching
Adjacent observed pairs	347,533 transitions	—	Date-gap audit
Exactly one-day adjacent pairs	196,611 transitions	—	Potential exact-day transitions before shared comparator filtering
Shared state cohort	87,944 transitions	31,700	Exact subsequent-calendar-day models

Table S2: **Developmental variables and their roles.**

Variable	Definition	Primary PCA status	Interpretation and limitation
Allocation bins 00–19	Twenty mass-conserving normalized growth shares over relative developmental time	Included; only fitted inputs	Shape representation; tied to interpolation resolution and observation support

Variable	Definition	Primary PCA status	Interpretation and limitation
Final area	Catalog event area, source-verified	Excluded; projected afterward	Endpoint scale
Duration	Catalog duration and dated elapsed time, reported distinctly	Excluded; projected/adjusted	Event length and observation opportunity
Observation count	Number of dated observed increments	Excluded; projected/adjusted	Observation depth
Observed daily peak	Maximum dated increment	Excluded; projected afterward	Daily maximum, not instantaneous rate
Mean growth	Reconstructed total divided by observed time definition	Excluded	Absolute scale summary
Slenderness	Legacy mixed duration–area descriptor	Excluded	Not developmental shape
Front-loaded fraction	Allocation in first half of relative development	Excluded; response/trait	Ordering-sensitive concentration
Late-growth fraction	Allocation in final quarter	Excluded; response/trait	Ordering-sensitive terminal allocation
Peak timing	Relative position of observed daily peak	Excluded; response/trait	Poorly resolved for short histories
Terminal taper	Final increment divided by observed daily peak	Excluded; response/trait	Remaining growth at observed end
Normalized entropy	$-\sum p \log p / \log n$	Excluded; response/trait	Allocation evenness; invariant to permutation
Detected pulses	Local maxima with at least 20% peak prominence	Excluded; response/landmark	Threshold-dependent episode count
Reactivation	Return above 25% peak after two subthreshold observed days	Excluded; response/landmark	Requires consecutive observation support
Normalized first/second differences	Differences between adjacent allocation bins	Excluded; responses/traits	Shape roughness, not physical velocity or acceleration
Weather variables	Daily gridMET summaries	Excluded; external predictors	Retrospective gridded exposure

Table S3: **PCA variance and stability across observation thresholds, evaluated using the same reference fires.**

Minimum consecutive observations	N	PC1 variance	First-five variance	Axis 1 correlation	Axis 2 correlation	Distance Spearman	15-neighbor overlap	Exemplar-tail Jaccard
2	27,499	0.51354	0.92466	0.98701	0.99726	0.99881	0.84613	0.66795
3, primary	10,246	0.34150	0.89444	1.00000	1.00000	1.00000	1.00000	1.00000
5	2,887	0.26621	0.82432	0.98709	0.99314	0.98767	0.77020	0.47053
7	1,171	0.23628	0.74214	0.98487	0.99664	0.96891	0.71787	0.60597

Table S4: **Temporal-order references and endpoint projections.**

Analysis	Quantity	Observed or PC1	Reference or PC2	Interpretation
Order shuffle	First-half allocation	0.54063	0.49958	Observed histories are more front loaded
Order shuffle	Detected pulses	1.24900	1.41280	Observed ordering yields fewer detected episodes
Order shuffle	Reactivation count	0.03800	0.12031	Observed ordering yields fewer reactivations
Order shuffle	First-five variance coverage	0.89582	0.85545	Shuffling changes compression
Dirichlet reference	First-five variance coverage	0.89582	0.90158 ($\alpha = 1$)	Positive allocations also compress strongly
Dirichlet reference	First-five variance coverage	0.89582	0.90853 ($\alpha = 10$)	Low dimensionality is not unique to observed histories
Endpoint projection	Final area	−0.308	0.259	Scale is excluded but associated with shape
Endpoint projection	Duration/count	−0.185	0.256	Length is excluded but associated with shape
Endpoint projection	Observed daily peak	−0.332	—	Daily growth magnitude is excluded but associated with PC1

Table S5: **Selected adjusted weather associations and response-specific predictive performance.**

Response or association	Estimate	Region R^2	Space–time R^2	Year R^2	Random R^2	Interpretation
Adjusted rank associations						
Mean VPD vs front loading	partial $\rho = 0.04989$	—	—	—	—	Small adjusted association
Mean VPD vs late allocation	partial $\rho = -0.06936$	—	—	—	—	Strongest highlighted adjusted association remains small
Mean precipitation vs late allocation	partial $\rho = 0.04072$	—	—	—	—	Small adjusted association
Mean precipitation vs pulses	partial $\rho = 0.02492$	—	—	—	—	Region interval crosses zero
Predictive performance: core weather means plus maximum VPD						
Front-loaded fraction	—	0.0215	0.0210	0.0356	0.0376	Weak predictive skill
Late-growth fraction	—	0.0383	0.0370	0.0500	0.0516	Weak predictive skill
Pulse count	—	0.0191	0.0183	0.0204	0.0213	Weak predictive skill
Peak timing	—	0.0005	0.0003	0.0021	0.0024	Essentially unrecovered
Reactivation count	—	0.0220	0.0208	0.0247	0.0270	Weak skill
Normalized entropy	—	0.0744	0.0725	0.1024	0.1043	Greater but limited skill
Terminal taper	—	0.0803	0.0784	0.0898	0.0918	Greater but limited skill
Normalized first difference	—	0.1755	0.1704	0.1971	0.2022	Highest displayed shape-trait skill

Response or association	Estimate	Region R^2	Space–time R^2	Year R^2	Random R^2	Interpretation
Normalized second difference	—	0.1606	0.1565	0.1780	0.1823	Response-specific roughness skill
Shape PC1, fold fit	—	0.0066	0.0029	0.0386	0.0415	Weak transfer of fold-specific coordinate

Table S6: **Exact-day state-model skill and increments above autoregressive state.**

Validation	Predictor set	Held-out R^2	ΔR^2 vs state	VPD increment beyond other interactions	Scope
Random fire	Autoregressive	0.45838	0	—	Interpolation diagnostic
Random fire	State + weather	0.46376	0.00539	—	Small additive weather gain
Random fire	All interactions	0.47359	0.01522	0.00632	Conditional associations
Year block	Autoregressive	0.45822	0	—	Temporal transfer
Year block	State + weather	0.46333	0.00511	—	Small additive weather gain
Year block	All interactions	0.47288	0.01465	0.00653	Conditional associations
Region block	Autoregressive	0.44764	0	—	Geographic transfer
Region block	State + weather	0.45265	0.00501	—	Small additive weather gain
Region block	All interactions	0.46579	0.01815	0.01179	Conditional associations
Space–time	Autoregressive	0.44753	0	—	Strict crossed transfer
Space–time	State + weather	0.45246	0.00493	—	Small additive weather gain
Space–time	All interactions	0.46553	0.01800	0.01199	Conditional associations

Table S7: **Matching coverage and mismatch relative to conditional permutations.**

Quantity	Weather match	Morphology match	Weather null	Morphology null
Unique pairs	3,710	3,145	—	—
Fraction of eligible fires paired	0.805	0.683	—	—
Ten-neighbor candidate coverage	0.912	0.813	—	—
Median other-space distance	0.99750	0.86719	1.00907	0.86978
Fraction other-space distance > 1	0.49730	0.39491	0.50401	0.40227
Null central interval for fraction > 1	—	—	0.49204–0.51834	0.39014–0.41750
Upper-tail probability	—	—	0.84158	0.83168

Table S8: **Selection into the primary gap-free developmental population.** Counts aggregate the validated inclusion–exclusion table across weather-completeness status. Percentages describe representation in the 10,246-fire primary cohort, not sampling weights.

Archive stratum	Archive events	Primary events	Retained (%)
Broad region			
Alaska	41,279	1,028	2.49
Central	89,584	2,339	2.61
East	43,972	1,418	3.22
Hawaii	55	6	10.91
Intermountain	58,896	3,259	5.53
West	44,783	2,196	4.90
Final area (km^2)			
≤ 1	205,970	1,799	0.87
> 1–10	64,297	6,830	10.62
> 10–100	7,531	1,384	18.38
> 100–1000	735	224	30.48
> 1000	36	9	25.00
Recorded duration (days)			

Archive stratum	Archive events	Primary events	Retained (%)
≤ 1	161,073	0	0.00
> 1–2	17,253	0	0.00
> 2–3	18,513	5,123	27.67
> 3–5	30,148	3,320	11.01
> 5–10	36,288	1,482	4.08
> 10–30	15,002	312	2.08
> 30	292	9	3.08

Table S9: **Primary standardized-Euclidean and Hellinger morphospace comparison.**

Metric	Primary	Hellinger or aligned value	Interpretation
PC1 variance explained	0.34150	0.33791	Similar leading compression
PC1–PC5 cumulative variance	0.89444	0.88282	Similar five-axis coverage
Aligned PC1 score correlation	—	0.94197	Strong axis correspondence
Aligned PC2 score correlation	—	0.93833	Strong axis correspondence
Five-dimensional Procrustes similarity	—	0.98479	Global configuration preserved
Five-dimensional score-subspace overlap	—	0.96009	Leading score spaces overlap strongly
Pairwise-distance Spearman correlation	—	0.97629	Global distance ranking preserved
Mean 15-neighbor overlap	—	0.70140	Local membership less invariant
Exemplar-tail Jaccard overlap	—	0.43027	Extreme membership metric dependent
PC1 correlation with front loading	−0.96834	−0.94502	Early-versus-late gradient persists

Supplementary References

References

[1] L. Giglio, L. Boschetti, D. P. Roy, M. L. Humber, C. O. Justice, The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sens. Environ.* **217**, 72–85 (2018).

[2] J. K. Balch, L. A. St. Denis, A. L. Mahood, N. P. Mietkiewicz, T. M. Williams, J. McGlinchy, M. C. Cook, FIRED (Fire Events Delineation): An open, flexible algorithm and database of US fire events derived from the MODIS burned area product (2001–2019). *Remote Sens.* **12**, 3498 (2020).

[3] A. L. Mahood, E. J. Lindrooth, M. C. Cook, J. K. Balch, Country-level fire perimeter datasets (2001–2021). *Sci. Data* **9**, 458 (2022).

[4] C. E. Shannon, A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).

[5] I. T. Jolliffe, *Principal Component Analysis*, ed. 2 (Springer, 2002).

[6] J. T. Abatzoglou, Development of gridded surface meteorological data for ecological applications and modelling. *Int. J. Climatol.* **33**, 121–131 (2013).